The Shirley reduced basis: a reduced order model for plane-wave DFT

Maxwell Hutchinson
*The Physics Department, University of Chicago, Chicago IL 60637**

David Prendergast
The Molecular Foundry, Lawrence Berkeley National Laboratory, Berkeley CA 94720
(Dated: March 3, 2014)

The Shirley reduced basis (SRB) represents the periodic parts of Bloch functions as linear combinations of eigenvectors taken from a coarse sample of the Brillouin zone, orthogonalized and reduced through proper orthogonal decomposition. We describe a novel transformation of the self-consistent density functional theory eigenproblem from a plane-wave basis with ultra-soft pseudopotentials to the SRB that is independent of the k-point. In particular, the number of operations over the space of plane-waves is independent of the number of k-points. The parameter space of the transformation is explored and suitable defaults are proposed. The SRB is shown to converge to the plane-wave solution. For reduced dimensional systems, reductions in computational cost, compared to the plane-wave calculations, exceed 5x. Performance on bulk systems improves by 1.52x in molecular dynamics-like contexts. This robust technique is well-suited to efficient study of systems with stringent requirements on numerical accuracy related to subtle details in the electronic band structure, such as topological insulators, Dirac semi-metals, metal surfaces and nanostructures, and charge transfer at interfaces with any of these systems. The techniques used to achieve a k-independent transformation could be applied to other computationally expensive matrix elements, such as those found in density functional perturbation theory and many-body perturbation theory.

PACS numbers: 71.15.-m, 71.15.Ap, 71.15.Dx, 02.70.-c, 02.60.-x

arXiv:1402.7366v1 [cond-mat.mtrl-sci] 28 Feb 2014

I. INTRODUCTION

Electronic structure is a cornerstone of modern materials science and condensed matter research efforts. Rapid advancements in experimental techniques have put pressure on the simulation community to provide efficient access to chemically accurate numerical predictions for a diverse range of materials, while high-throughput efforts [1] have embraced standard electronic structure methods, such as density functional theory (DFT)[2], to explore entire databases of materials[3]. For materials science applications, typically focused on condensed phases, surfaces and extended nanostructures, the use of atomistic models employing periodic boundary conditions is the norm, with plane-wave DFT (PWDFT) providing a robust numerical approach for systems where the number of atoms does not extend to thousands. While PWDFT has been shown to be adequately robust in the treatment of exotic materials, whose electronic properties are defined by subtle details in electronic band structure (e.g., Dirac semi-metals and topological insulators [4]), necessarily numerically converged calculations of complex systems often come at prohibitive computational cost. The Shirley reduced basis (SRB) technique is able to substantially reduce the numerical effort of PWDFT while maintaining uniform basis set convergence.

PWDFT is a natural numerical representation for electronic structure under periodic boundary conditions, which are commonly employed to model crystalline and amorphous condensed phases, surfaces and their adsorbates, materials interfaces or heterojunctions, and extended nanostructures, such as sheets, ribbons, wires and tubes. PWDFT can also be employed to study finite or molecular systems, but that will not be the focus of this work. Note that finite temperature first-principles molecular dynamics simulations for condensed phases, surfaces, interfaces and extended nano structures typically also often rely on PWDFT, due to the robust description of electronic structure provided by this numerical representation, even for configurations far from equilibrium. In particular, we are interested in systems which are expensive to model due to the necessity to capture subtle details of electronic band structure which define the function of these materials. Typically, this complexity is related to adequately describing the Fermi surface of complex metals or semi-metals, which requires detailed knowledge of electronic band structure with respect to electron wave vectors.

Bloch's theorem factors the representation of wave functions in periodic systems into a slowly oscillating phase, defined by the electron wave vector k, and a strictly periodic function, $u(r)$[5]:

$$\psi(r) = e^{ik\cdot r}u(r) \qquad u(r+a) = u(r), \tag{1}$$

where a is a Bravais lattice vector of the infinite periodic system. In DFT, the Hamiltonian is block-diagonal with respect to the k-point, however, the expectation value of the electron density relies on integrating contributions over all values of k in the first Brillouin zone (BZ). Numerical solutions to electronic structure problems discretely sample the BZ with a finite set of k-points, $k \in K \subset BZ$. Existing electronic structure methods

* maxhutch@uchicago.edu

treat the k-dependent eigen-problems independently. In reality, the eigenproblems and associated eigensolutions at each k-point contain a degree of redundancy. The periodic Bloch states, $|u_{nk}\rangle$ are known to have relatively weak k-dependence in comparison to the dispersion relation ε_{nk}, which is mainly driven by the quadratic k-dependence of the kinetic energy.

Shirley proposed that a reduced basis comprised of Bloch states over a coarse sample of the BZ be used to represent the eigenproblem throughout the BZ [6]. We denote the initial coarse sample as the *q-points*. Thus, the nth Bloch state at point k can be represented as

$$|u_{nk}\rangle = \sum_{mq} |u_{mq}\rangle \, c_{mq}^{nk}. \tag{2}$$

The number of q-points needed to accurately reproduce the solutions at all k-points is small. We refer to this technique, the use of a coarse BZ sample as a basis, as *Shirley interpolation*. Analogies can be drawn to $k \cdot p$ theory[7, 8], which can effectively interpolate electronic structure in the close neighborhood of a specific k-point using a large eigenspectrum from that k-point alone. By contrast, Shirley interpolation aims to provide a single reduced basis which spans the entire BZ by combining details from multiple q-points and a relatively small number of bands.

The basis size can be further reduced by selecting linear combinations of q-point states though a principal component analysis [6]. The principal components with small eigenvalues are removed from the basis, provided the specification of a tolerated error. The use of principal component analysis in reducing the dimensionality of PDEs is often called proper orthogonal decomposition (POD) [9].

The *Shirley reduced basis* (SRB) is the union of Shirley interpolation and proper orthogonal decomposition. Shirley applied the method to the non-self-consistent evaluation of band structure and associated spectroscopy and achieved speedups in excess of 1000× for dense k-point samples [6].

Prendergast and Louie extended Shirley's work to ultra-soft pseudopotentials (USPP) through interpolation of the projectors across the BZ [10]. They also recognized that preserving periodicity with respect to $\mathbf{k}$ across the first BZ could be achieved using periodic images of the original coarse q-point basis. In particular, the periodic images of the Γ-point were used to successfully enforce the periodic symmetry of the band structure. Such images can be generated merely via mathematical transformation, without additional expensive plane-wave calculations.

Here, we further extend the method to self-consistent calculations, employing a new set of techniques to mitigate the added cost of constructing an electron density [2]. We also remove the interpolation of the atomic projectors required for the non-local potential, in favor of an auxiliary basis approach similar to the SRB for the periodic component of the Bloch eigenfunctions. We implement the SRB in the popular PWDFT code Quantum Espresso [11]. Our implementation is available for general use [12].

The approach is shown to be uniformly convergent and numerically efficient, achieving significant speedups for self-consistent calculations, including relaxations. The accurate reproduction of forces and stresses also indicate the possibility for significant speed up of first-principles molecular dynamics simulations, which can prove prohibitively expensive for systems which require k-point sampling to capture metallic or semi-metallic behavior or the possibility of a transition to a metallic state via phase change [13] or charge transfer [14]. One can view such molecular dynamics simulations simplistically as requiring a converged self-consistent-field calculation for each time step and determination of the next atomic configuration based on the computed forces and stresses. Therefore, significant reductions in computing time can be realized even for modest improvements in the time to accurately compute a single step.

In Section III, we introduce the self-consistent extension of Shirley's method, the SRB, in general terms. Then, in Section V, we describe the parameters of the method and their affect on the solution. In Section VI, we demonstrate the convergence properties of SRB. In Section VII, we present a performance model for the SRB. In Section II, we discuss related work on reduced order models in other fields. In Section X, we conclude by recognizing a class of problems that are well suited for the SRB and describing a way forward for reduced order models for electronic structure calculations.

II. RELATED WORK

The SRB is a reduced order method (ROM) based on Shirley interpolation and proper orthogonal decomposition (POD), which is related to principal component analysis and Galerkin projection in other fields [15]. POD, principal component analysis, and Galerkin projection have been well studied in fields as diverse as fluid dynamics [16], machine learning[17], and image processing[9, 18]. Shirley interpolation appears to have originated with the work of Shirley himself.

ROMs were developed in the fluid dynamics community [16, 19] as a means to reduce the dimensionality of steady-state turbulent flows. Samples of the velocity fluctuations $u - \langle u \rangle$ taken at intervals longer than the correlation time are fed into a singular-value decomposition to identify coherent structures in the flow. The governing equations are then expressed in terms of the coherent structures, creating a model for the time-dependent energy transfer between the structures. The SRB differs from common ROMs in two ways: 1) it models a non-linear system and 2) it samples in Bloch space. The non-linearity makes a-priori error estimates difficult. The Bloch-space sampling provides structure to the operator transformations.

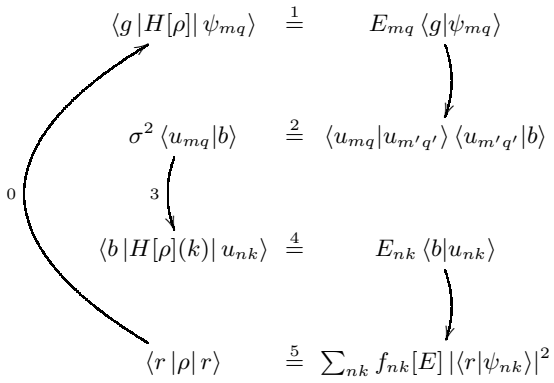

FIG. 1: Flow chart of the SRB method. The regular plane-wave SCF method is embodied in steps 0, 1, and 5. The SRB adds step 2 to build the reduced basis, step 3 to transform the k-dependent Hamiltonian into the reduced basis, and step 4 to solve the reduced eigenproblem.

The *optimal representation of the polarization propogator* in [20] uses the SVD of the product basis of Wannier-function expansions of wave functions to reduce the basis size, akin to POD. The representation does not, however, take advantage of the smoothness of a solution manifold as the SRB does in k-space and fluid dynamics ROMs do in time. In other words, it lacks the interpolation step. The optimal representation is most like the auxiliary basis used in this work, Section III C, for the transformation of projectors and the reduced rank density matrix, neither of which involve k-space interpolation.

The $k\cdot p$-method has been applied to density functional theory calculations by Persson and Ambrosch-Draxl[21]. It relies on analytic k-dependence of the Hamiltonian to express linear relations between the Bloch states at different k-points. That reliance, however, restricts $k \cdot p$ to application to methods with only local potentials, such as FLPAW. The SRB has no such restrictions, and is presented here with the generality of ultra-soft pseudopotentials.

Reduced Bloch Mode Expansion (RBME) is an expansion in Bloch eigenfunctions nearly identical to the Shirley interpolation [22] step, but lacking the POD step. It has wide application beyond electronic structure to band structures of periodic systems and is agnostic to the discretization of the Bloch states. The RBME technique post-dates Shirley's original work.

The "reduced-basis" method described by Pau[23] makes similar assumptions about the form of the potential in order to produce an eigenproblem that is affine in k. Like RBME, it lacks the POD step.

III. SELF-CONSISTENT ALGORITHM

The methods of Shirley, Prendergast, and Louie for non-self-consistent interpolation are well described in the literature [6, 10]. The most basic idea is to use the span of the periodic components of the Bloch eigenfunctions at a small number of points in the first BZ as a basis for the entire BZ. Furthermore, a POD of these periodic functions allows for the basis size to be reduced while preserving the trace of physical operators within a given tolerance. Finally, a k-dependent Hamiltonian can be constructed without an explicit change of basis through polynomial expansion of the kinetic energy and interpolation of projectors of the non-local potential.

The algorithm presented below is a natural generalization to the self-consistent regime. It differs primarily in the addition of a step to reconstruct a real-space electron density from Bloch eigenfunctions in the reduced basis, and secondarily in more general approaches to the basis transformations and interpolations.

A few notes: We use the term *expansion* to refer to exact polynomial forms and reserve *interpolation* for methods that include truncation errors. The basis produced by this scheme is referred to as 'reduced'[15], as opposed to 'optimal' in the previous literature, as it is only optimal within the space spanned by the periodic components of the input Bloch eigenfunctions. For the sake of clarity, we assume the conventional method employs a plane-wave basis with pseudopotentials [24, 25], though the method generalizes to any discretization of the Bloch eigen-functions, for example, in real-space using analytic or numerical atom-centered orbitals or using projector-augmented waves [26]. We use the unit cube in reciprocal lattice coordinates as our BZ, as opposed to the conventional choice of the origin-centered Wigner-Seitz cell.

The algorithm refines an electron density. Generally, it would be repeated until a convergence criterion is met, signaling self-consistency. The outline of the algorithm can be seen in Figure 1: Start with an electron density represented in real-space. 0) Represent the Hamiltonian as a functional of the electron density using a conventional basis, $\{|g_i\rangle\}$; 1) Compute Bloch eigenfunctions $\langle g|\bar{\psi}_{nq}\rangle$ using a conventional method over a coarse reciprocal mesh, Q; 2) Create a reduced basis $\{|b_i\rangle\}$ from an optimal subspace of the span of the periodic components of the conventional Bloch eigenfunctions; 3) Represent k-dependent Hamiltonians in the new basis over a fine reciprocal mesh K; 4) Solve the k-dependent Hamiltonians in the new basis producing the periodic components of the eigenfunctions $\langle b|u_{nk}\rangle$; 5) Construct a real-space electron density from the eigenfunctions in the new basis.

A. Constructing the basis

The conventional algorithm produces eigenfunctions which are Bloch functions $|\bar{\psi}_{mq}\rangle = e^{iq\cdot r}\,|\bar{u}_{mq}\rangle$ on a reciprocal mesh Q. Symmetries in the problem can be used

to map the Bloch eigenfunctions computed at a single q-point to other points in reciprocal space. In addition, the periodicity of eigensolutions in k-space permits states from the surface of the first BZ, e.g. the reciprocal space origin $q = \Gamma = 0$, to be translated by any other reciprocal lattice vector, e.g. the corners of the unit cube, $0^3 \to \{0,1\}^3$ in reciprocal lattice coordinates. This procedure is further described in [10], and generally requires little numerical effort, given that such translations can be achieved merely by a reordering of the Fourier coefficients of each state.

Additional symmetries can provide more input q-points without additional plane-wave computation. For example, inversion symmetry is computed by conjugating the plane-wave coefficients:

$$c_{n,q} = c^*_{n,-q}.$$

which provides two interior k-points for the cost of one. The symmetries can be compounded. For example, a single q-point on the edge can be used to construct 7 other q-points without significant computation assuming inversion symmetry and periodicity:

$$\begin{aligned}(0,0,1/4) \to \{&(0,1,1/4),(1,0,1/4),(1,1,1/4),\\ &(0,0,3/4),(0,1,3/4),(1,0,3/4),(1,1,3/4)\}\end{aligned}$$

We use POD to pick a subspace of the span of the periodic components of the Bloch eigenfunctions that is optimally representative. First, construct a covariance (or overlap) matrix:

$$C_{i,j} \equiv \langle u_i | u_j \rangle \,, \tag{3}$$

where i, j are composite indices over m, q, the space of coarsely sampled eigen-functions. Next, diagonalize the covariance matrix:

$$\sum_\gamma \langle u_j | u_\gamma \rangle \langle u_\gamma | b_i \rangle = \sigma_i^2 \langle u_j | b_i \rangle \,, \tag{4}$$

where the eigenvalues σ_i^2 are the variances captured by the basis elements $|b_i\rangle$. Select those basis elements, $|b_i\rangle$, with the largest variances as the reduced basis. The size of the basis can be informed by choosing a sufficient number of basis elements to capture a significant fraction of the total variance. For example, the missing variance could be constrained to be below a certain fraction of the total variance:

$$1 - \frac{\sum_i^{N_b} \sigma_i^2}{\sum_j^N \sigma_j^2} < \sigma_b^2, \tag{5}$$

where N_b is the number of basis elements, N is the rank of the overlap matrix C, and σ_b^2 is the maximum tolerated error.

The basis elements are represented with respect to the original periodic components of the Bloch eigenfunctions, so one must generally transform them back to the original basis. In the case of plane-waves, the transformation takes the form:

$$\langle g_j | b_i \rangle = \sum_\gamma \langle g_j | u_\gamma \rangle \langle u_\gamma | b_i \rangle \,. \tag{6}$$

B. k-dependent Hamiltonians

Armed with this reduced basis, we use it to expand a k-dependent Hamiltonian $H(k) = e^{-ik\cdot r} H e^{ik\cdot r}$. In the plane-wave pseudopotential framework, the Hamiltonian has three components: the kinetic energy, the local potential, and the non-local potential [5, 27]:

$$H \equiv |g\rangle K \langle g| + |r\rangle V \langle r| + \left|\beta\right\rangle V^{nl} \left\langle\beta'\right| . \tag{7}$$

It was previously shown that the kinetic energy could be written as a quadratic expansion in $\vec{k}$:

$$2\left\langle g_i \left| K(\vec{k}) \right| g_i \right\rangle = \left|\vec{k} + \vec{g}_i\right|^2 = |\vec{k}|^2 + 2\vec{g}_i \cdot \vec{k} + |\vec{g}_i|^2, \tag{8}$$

which can be transformed to the reduced basis by means of the matrices:

$$\left\langle b_i \left| K^0 \right| b_j \right\rangle = \sum_\gamma \langle b_i | g_\gamma \rangle \, |\vec{g}_\gamma|^2 \, \langle g_\gamma | b_j \rangle \tag{9}$$

$$\left\langle b_i \left| \vec{K}^1 \right| b_j \right\rangle = \sum_\gamma \langle b_i | g_\gamma \rangle \, \vec{g} \, \langle g_\gamma | b_j \rangle \tag{10}$$

$$\left\langle b_i \left| K^2 \right| b_j \right\rangle = \delta_{ij} = 1, \tag{11}$$

such that

$$2\left\langle b_i \left| K(\vec{k}) \right| b_j \right\rangle = |\vec{k}|^2 K^2 + 2\vec{k} \cdot \vec{K}^1 + K^0, \tag{12}$$

which exactly isolates the k dependence in the transformation of the kinetic energy.

The local potential is k-independent. Its reduced basis representation is computed by transformation from the plane-wave basis to real-space:

$$\begin{aligned}&\left\langle b_i \left| V^{loc} \right| b_j \right\rangle = \\ &\qquad \sum_{\gamma,\gamma',r} \langle b_i | g_\gamma \rangle \, \langle g_\gamma | r \rangle \, V^{loc}(r) \, \langle r | g_{\gamma'} \rangle \, \langle g_{\gamma'} | b_j \rangle \,,\end{aligned} \tag{13}$$

where the sums over γ, γ' can be implemented as Fourier transformations.

The non-local potential is generally expressed as a matrix in the basis of atomic pseudo-wave functions, or just pseudo-functions: $\left\langle \tilde{\phi}_{a,l'} \left| V^{nl} \right| \tilde{\phi}_{a,l} \right\rangle$, where a runs over atomic sites and l over composite angular momentum quantum numbers [28]. The non-local potential is diagonal for NCPP and block diagonal with respect to the atomic index for USPP. We will occasionally compress the indices $(a, l) \to \alpha$. The space of atomic wave functions is accessed through a dual-space of projectors, β,

which couples isolated atomic Hilbert spaces to the overall periodic space comprising multiple atomic sites:

$$\sum_{l,l'} \left|\beta_{a,l}\right\rangle \left\langle\tilde{\phi}_{a,l'}\right| = I_a, \tag{14}$$

where I_a represents the identity in the space spanned by the pseudo-functions in a finite volume around atomic site a. These matrix elements do not depend on k, but the projectors $|\beta_{a,l}\rangle$ do. Transforming the projectors is all that must be done to transform the non-local potential to the reduced basis:

$$\langle b_i|\beta_\alpha(k)\rangle = \sum_\gamma \langle b_i|g_\gamma\rangle \langle g_\gamma|\beta_\alpha(k)\rangle \tag{15}$$

$$\left\langle b_i \left|V^{nl}(k)\right| b_j\right\rangle = \sum_{\alpha,\alpha'} \langle b_i|\beta_\alpha(k)\rangle \left\langle\tilde{\phi}_\alpha \left|V^{nl}\right| \tilde{\phi}_{\alpha'}\right\rangle \langle\beta_{\alpha'}(k)|b_j\rangle . \tag{16}$$

For ultra-soft pseudo-potentials, the overlap matrix, S, is computed in the same way:

$$\langle b_i \left|S(k)\right| b_j\rangle - \delta_{i,j} = \sum_{\alpha,\alpha'} \langle b_i|\beta_\alpha(k)\rangle \left\langle\tilde{\phi}_\alpha \left|S\right| \tilde{\phi}_{\alpha'}\right\rangle \langle\beta_{\alpha'}(k)|b_j\rangle . \tag{17}$$

Note that when expressed in the basis of pseudo-functions, the non-trivial part of the overlap matrix, $\left\langle\tilde{\phi}_\alpha \left|S\right| \tilde{\phi}_{\alpha'}\right\rangle$, and non-local potential, $\left\langle\tilde{\phi}_\alpha \left|V^{nl}\right| \tilde{\phi}_{\alpha'}\right\rangle$, are often referred to as Q and D, respectively [24].

C. Auxiliary basis for projectors

The projectors can be written as the product of origin-centered, k-dependent atomic projectors, $\beta_{0,l}(k)$, and an atom-dependent structure factor, $\mathcal{S}(a)$:

$$\langle g_i|\beta_{a,l}(k)\rangle = \langle g_i \left|\mathcal{S}(a)\right| g_i\rangle \langle g_i|\beta_{0,l}(k)\rangle , \tag{18}$$

where l runs over angular momenta, a runs over the atomic centers, and $a = 0$ corresponds to the origin. The structure factor can be written as:

$$\langle g \left|S(a)\right| g\rangle = e^{-i(g+k)\cdot\tau_a}, \tag{19}$$

If the number of k-points is large, we could consider transforming the structure factor in the SRB before applying it to the origin-centered projectors:

$$\begin{aligned} \langle b_i \left|\mathcal{S}(a)\right| g_j\rangle &= \langle b_i|g_j\rangle \langle g_j \left|\mathcal{S}(a)\right| g_j\rangle \\ \langle b_i|\beta_{a,l}(k)\rangle &= \sum_\gamma \langle b_i \left|\mathcal{S}(a)\right| g_\gamma\rangle \langle g_\gamma|\beta_{0,l}(k)\rangle . \end{aligned} \tag{20}$$

The inner product in the plane-wave basis is inefficient: the size of the plane-wave basis is much larger than the rank of $\langle b_i \left|\mathcal{S}(a)\right| g_j\rangle$, which is the lesser of the size of the SRB and the dimension of the space spanned by the origin-centered projectors $\langle g_\gamma|\beta_{a,l}(k)\rangle$. We can introduce an auxiliary basis, $|x\rangle$, to reduce the effort in the inner product:

$$\begin{aligned} \langle b_i \left|\mathcal{S}(a)\right| x_j\rangle &= \sum_\gamma \langle b_i|g_\gamma\rangle \langle g_\gamma \left|\mathcal{S}(a)\right| g_\gamma\rangle \langle g_\gamma|x_j\rangle \\ \langle x_i|\beta_{0,l}(k)\rangle &= \sum_\gamma \langle x_i|g_\gamma\rangle \langle g_\gamma|\beta_{0,l}(k)\rangle \\ \langle b_i|\beta_{a,l}(k)\rangle &= \sum_j \langle b_i \left|\mathcal{S}(a)\right| x_j\rangle \langle x_j|\beta_{0,l}(k)\rangle . \end{aligned} \tag{21}$$

The basis x needs to represent the origin-centered projectors through the BZ. This is analogous to the SRB, which must represent the periodic Bloch functions through the BZ. Indeed, we use the procedure outlined in Section III A, but with $\langle g_i|\beta_l(q,0)\rangle$ as input states. This auxiliary basis need only be produced once, as it is independent of the electron density and atomic positions. Indeed, it could even be pre-computed and packaged with the pseudo-potential. Therefore, we prefer to use the entire K grid as the 'coarse' sample Q, which makes the fractional variance a good metric for the accuracy of the basis. In that sense, the auxiliary basis $\{|x_i\rangle\}$ is optimal. This procedure must be performed for each atomic species.

D. Diagonalizing the k-dependent Hamiltonian

The Hamiltonian and, in the case of USPP, the overlap matrix are dense. Therefore, it is natural to leverage highly optimized direct solvers, such as those found in the LAPACK library [29]. However, the ratio of the number of basis elements to the number of bands is $O(10)$, which is large enough to justify the use of an iterative solver. Furthermore, as in the conventional case, diagonalization in early SCF steps need not be fully converged. We intend to explore the use of iterative solvers in the future.

E. Constructing the density matrix

After diagonalizing the k-dependent Hamiltonian to get the periodic components of the eigenfunctions $\langle b_i|u_{mk}\rangle$ and the energies ε_{mk}, one must produce a real-space electron density so the process can be repeated. The simplest way of doing this would be to transform the eigenfunctions back to plane-waves and then accumulate the density in the usual way:

$$\begin{aligned} \langle g_i|u_{nk}\rangle &= \sum_j \langle g_i|b_j\rangle \langle b_j|u_{nk}\rangle \\ \langle r|u_{nk}\rangle &= \sum_j \langle r|g_j\rangle \langle g_j|u_{nk}\rangle \\ \langle r \left|\rho\right| r\rangle &= \sum_{nk} \langle r|u_{nk}\rangle f(\varepsilon_{nk}) \langle u_{nk}|r\rangle , \end{aligned} \tag{22}$$

where $f(\varepsilon)$ is some occupation function and the transformation from plane-waves to real-space is generally accomplished with an FFT.

Another method takes advantage of the small size of the SRB to form a density matrix:

$$\langle b_i \left|\rho\right| b_j \rangle = \sum_{nk} \langle b_i | u_{nk} \rangle \, f(\varepsilon_{nk}) \, \langle u_{nk} | b_j \rangle \,. \tag{23}$$

Using the Hermitian singular value decomposition $\rho = V \Sigma V^\dagger$, one can re-write the density matrix in terms of its rank-1 components:

$$\langle b_i \left|\rho\right| b_j \rangle = \sum_{\nu} \langle b_i | v_\nu \rangle \, \sigma_\nu \, \langle v_\nu | b_j \rangle \,, \tag{24}$$

where σ_ν are the diagonal elements of Σ. The rank-1 components, $\langle b_i | v_\nu \rangle$, can be accumulated in real-space just as we did for the periodic parts of the Bloch eigenfunctions in Equation 22:

$$\begin{aligned} \langle g_i | v_\nu \rangle &= \sum_j \langle g_i | b_j \rangle \, \langle b_j | v_\nu \rangle \\ \langle r | v_\nu \rangle &= \sum_j \langle r | g_j \rangle \, \langle g_j | v_\nu \rangle \\ \langle r \left|\rho\right| r \rangle &= \sum_\nu \langle r | v_\nu \rangle \, \sigma_\nu \, \langle v_\nu | r \rangle \,. \end{aligned} \tag{25}$$

This method has the advantage that the number of rank-1 components is bounded above by the size of the SRB. If there are more bands across all k-points than basis elements, $N_b N_k > N_r$, this approach performs fewer transformations. In practice, many of the σ_i are very small so they can be truncated just as small occupations $f(\varepsilon)$ generally are.

In the case of USPP, there is an addition to the electron density:

$$\begin{aligned} &\langle r \left|\tilde{\rho}\right| r \rangle = \\ &\sum_{nk,\alpha,\alpha'} f(\varepsilon_{nk}) \, \langle u_{nk} | \beta_\alpha(k) \rangle \left\langle \tilde{\phi}_\alpha \left| S(r) \right| \tilde{\phi}'_\alpha \right\rangle \langle \beta'_\alpha(k) | u_{nk} \rangle \,, \end{aligned} \tag{26}$$

which requires the inner product of the wave functions and the projectors:

$$\langle \beta_i(k) | u_{nk} \rangle = \sum_j \langle \beta_i(k) | b_j \rangle \, \langle b_j | u_{nk} \rangle \,. \tag{27}$$

Note that the length of the inner product is the reduced basis size, representing a significant reduction in work over the plane-wave counterpart.

F. Basis saving

The self consistent procedure serves not only to relax the electron density to the ground state, but also to relax the subspace spanned by the SRB to include the true ground state. Therefore, we consider two convergences: the convergence of the electron density and the convergence of the SRB. To fully generalize the method, we consider these two convergences separately. Namely, we can refrain from updating the SRB at every SCF iteration, instead 'saving' the transformation matrix, kinetic energy, projectors, and factorized overlap matrix from the previous iteration. The only parts of the eigenproblem which must be recomputed are the explicitly density dependent operators: the local and, in the case of USPP, non-local potentials, Equations 13 and 16, respectively. Relaxing the basis more slowly than the electron density can slow the convergence of the electron density, but the avoided cost of computing a new basis and transforming density independent quantities can result in performance improvement.

The two convergences need not be terminated by the same error condition. In general, the solution is less sensitive to the basis as it is to the electron density. Thus, the basis can be frozen on a weaker convergence threshold than the electron density while preserving the accuracy of the solution. In fact, if the basis is not frozen near convergence, it can lead to slowly decaying oscillations in the charge density and basis elements, similar to so-called charge sloshing, delaying convergence.

G. Forces and stresses

The electron density produced by the SRB technique, Equation 25, is represented in the same basis as the plane-wave result. Therefore to compute forces and stresses, we need only address terms that depend on the wave-functions directly [30].

The forces depend on derivative operators nominally expressed in the plane-wave basis:

$$\mathcal{D}^\mu(k) \equiv \left\langle g_i \left| \frac{\partial}{\partial R^\mu} \right| g_j \right\rangle = (g_i^\mu + k^\mu) \delta_{i,j} \,, \tag{28}$$

which are used to compute the derivatives of the projection:

$$\frac{\partial}{\partial R_a^\mu} \langle \beta_{a,l}(k) | u_{nk} \rangle = \left[\langle \beta_{0,l}(k) \left| \mathcal{S}(a)^\dagger \mathcal{D}^\mu(k)^\dagger \right] \right| u_{nk} \rangle \,. \tag{29}$$

The matrices $\mathcal{S}$ and $\mathcal{D}$ are Hermitian and diagonal in the plane-wave basis, so we can drop the adjoints and commute them. We need to compute the quantity:

$$\langle b_i \left| \mathcal{S}(a) \mathcal{D}^\mu(k) \right| \beta_{0,l} \rangle \,, \tag{30}$$

so we can take inner products with the periodic Bloch states $\langle b_i | u_{nk} \rangle$ in the reduced basis.

This problem is analogous to that of the projectors, Equation 27, but with $\mathcal{D}^\mu \beta_{0,l}$ replacing $\beta_{0,l}$. Therefore, we can form an auxiliary basis just as in Equation 21. As before, the auxiliary basis is independent of the electron density and the atomic position, so it can be produced

once and stored for the duration of the calculation. For relaxations and molecular dynamics, this allows the cost of producing the basis to be amortized by many force evaluations. For self-consistent calculations with a single force evaluation, it may be advantageous to compute the derivatives of the projection in the plane-wave basis.

The stress has an additional term that depends on the wave-functions, the kinetic stress [31]:

$$\sigma_{\text{kin}}^{\mu,\nu} = \sum_{n,k,g} \langle u_{nk}|g\rangle \, (g+k)^{\mu}(g+k)^{\nu} \, \langle g|u_{nk}\rangle \,, \tag{31}$$

which leads to the k-dependent kinetic stress operator:

$$\langle g \, |\hat{\sigma}_{\text{kin}}^{\mu,\nu}(k)| \, g\rangle = (g+k)^{\mu}(g+k)^{\nu}. \tag{32}$$

Unsurprisingly, this term can be treated as the kinetic energy operator was in Equation 8: the product is expanded and the k-dependence factored yielding a k-independent transformation:

$$\langle g \, |\hat{\sigma}_{\text{kin}}^{\mu,\nu}(k)| \, g\rangle = g^{\mu}g^{\nu} + g^{\mu}k^{\nu} + g^{\nu}k^{\mu} + k^{\mu}k^{\nu}, \tag{33}$$

the following nine transformations are required:

$$\langle b_i \, |g^{\mu}g^{\nu}| \, b_j\rangle \qquad = \sum_{g} \langle b_i|g\rangle \, g^{\mu}g^{\nu} \, \langle g|b_j\rangle \tag{34}$$

$$\langle b_i \, |g^{\mu}| \, b_j\rangle \qquad = \sum_{g} \langle b_i|g\rangle \, g^{\mu} \, \langle g|b_j\rangle \,, \tag{35}$$

which are then composed as:

$$\begin{aligned}\langle b_i \, |\hat{\sigma}_{\text{kin}}^{\mu,\nu}(k)| \, b_j\rangle &= k^{\mu}k^{\nu}\delta_{i,j} \\ &+ \langle b_i \, |g^{\mu}g^{\nu}| \, b_j\rangle + \langle b_i \, |g^{\mu}| \, b_j\rangle \, k^{\nu} + \langle b_i \, |g^{\nu}| \, b_j\rangle \, k^{\mu}.\end{aligned} \tag{36}$$

Conveniently, the single-g terms were already computed for the kinetic energy in Equation 10. Further, one of the diagonal double-g terms can be computed by subtraction:

$$g^3g^3 = |g|^2 - g^1g^1 - g^2g^2, \tag{37}$$

with the $|g|^2$ term previously defined as K^0, Equation 9. Thus, the total number of additional diagonal transformations is five.

The second derivatives found in the non-local contribution to the stress can be computed in a similar way as Equation 30. There are twelve such terms but they are density and atomic position independent, so they can be re-used.

IV. EXAMPLES

Through the rest of this paper, we discuss results from three representative systems: a $(3,3)$ carbon nanotube (CNT), a nickel slab decorated with water and hydroxide, and a bulk gold molecular dynamics snapshot at 2000 K. We will explain in the following sections why they are representative in the context of the SRB. The raw input files used for these examples can be found in the `srb-supplemental` repository[32].

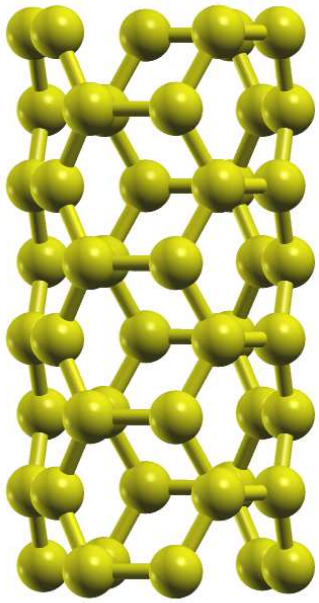

FIG. 2: Four cells of a (3,3) CNT.

A. (3,3) carbon nanotube

The (3,3) carbon nanotube (CNT) is a metallic one-dimensional nano-structure. It has a 12-atom primitive cell arranged around the z-axis. To replicate isolation within periodic boundary conditions, a large vacuum regions is added in the x-y plane. We are only concerned with relaxing stress along the tube axis.

Metallic systems with a small number of Fermi level crossings, such as the 'Dirac points' in CNTs, require thorough sampling of the BZ. When the BZ is undersampled, the crossings can be missed, leading to an artificial band gap. In this case, 32 irreducible points are required.

The CNT is representative of a large class of extended one-dimensional metallic and semiconducting molecules (e.g. linear polymers) and nanostructures, including nanotubes, nanowires, and nanoribbons. Accurate descriptions of the electronic structure of these materials are vital to estimating their electronic properties and their structures. In addition, the proximity of metallic or semimetallic nanostructures to molecules, nanoparticles or extended substrates often results in significant charge transfer across the interface due to electronic coupling and hybridization [33]. Inaccurate descriptions of the Fermi surface can modify the details of electron transfer and lead to poor descriptions of chemisorption or electronic transmittance across such interfaces. Modeling such interactions at finite temperature would also require sufficient k-point sampling to cover the worst possibilities during an entire trajectory – such as switching from metallic to semi-metallic behavior [34]. In the absence of numerical convergence, such simulations would incorrectly describe

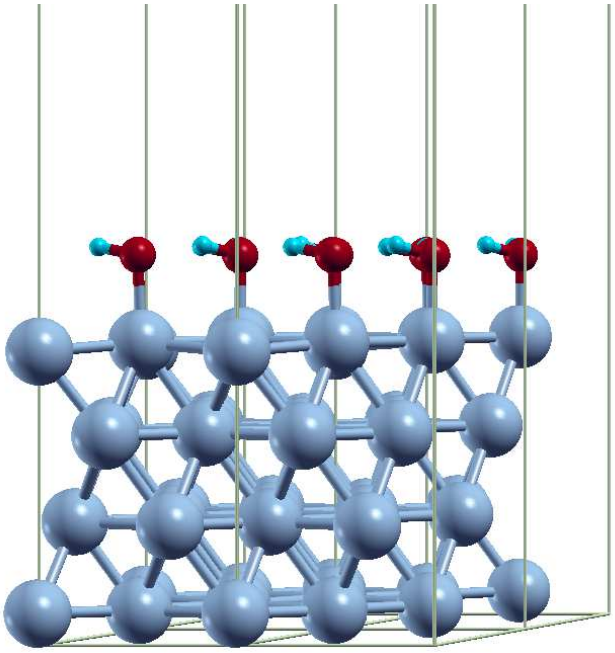

FIG. 3: Nickel slab decorated with water and hydroxide. Cell is repeated 2×2 in the plane of the slab.

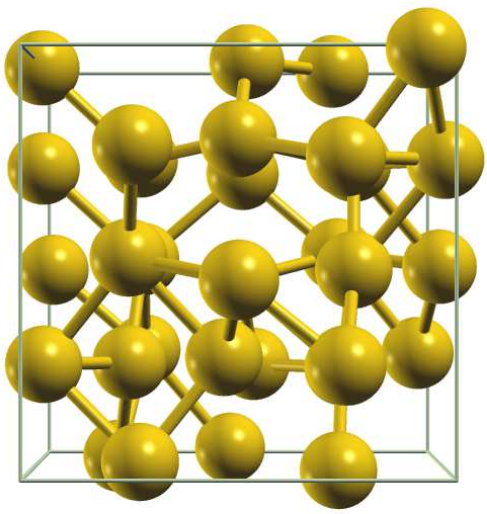

FIG. 4: Gold snapshot at 2000K. Cell contains 32 atoms.

the system as semiconducting in both cases.

B. Nickel slab

We provide an example relevant to the study of surface chemistry and catalysis. A four atomic layer Ni slab is chosen to model a semi-infinite reactive metal surface, with sufficient vacuum padding to reduce the influence of electronic coupling between the top and bottom surfaces. We decorate one side of the slab with water molecules and hydroxide moieties, representing a proposed surface coverage of Ni in the presence of water vapor [35]. The particular details of this system are not important within the context of the current study, and we will not try to draw any physical or chemical conclusions. However, this quasi-2D system is representative of a large class of simulations which aim to model surfaces and interfaces and their chemistry or reactivity. Specific details of surface relaxation and reconstruction, charge transfer and reactivity, and electronic coupling via hybridization are all strongly dependent on an accurate description of the electronic structure and require numerically converged sampling of the BZ. The combination of a large number of k-points and the additional requirement to include a large number of plane-waves to describe the vacuum above and below the slab render such common calculations prohibitively expensive. Any gains in efficiency through the use of the SRB would surely be welcome in the surface science and catalysis communities.

C. Gold snapshot

The gold system is a 32 atom disordered snapshot extracted from an MD simulation at 2000 K [36]. A hard norm-conserving pseudo-potential is used to describe the ionic cores under these extreme conditions. The gold snapshot is representative of the typical accuracy requirements and computational cost per step of first-principles molecular dynamics of metals under extreme conditions. We wish to demonstrate the possibility of realizing some speedup for such systems. In addition, we wish to provide a practicable option for running robust simulations of systems which undergo insulator to metal transitions or pressure induced metallization. Furthermore, we hope to enable additional parallelization paradigms for large scale molecular dynamics simulations through our use of the SRB.

D. Agreement

In the following sections, we will compare different calculations of the same physical system. Here, we define the standards by which we compare calculations. That is, what conditions need to be met for two calculations to be in 'agreement'.

We define agreement in terms of three measures: the band structure (ε), the force (F), and pressure (P). Agreement in band structure means less than 5meV absolute root mean square error over bands that are below or cross the Fermi level. We zero the band structure at

System	Calculation	K	R_v	A	PP	E_{xc}
CNT	SCF/Bands	33	$60a_0$	12	US	PBE
Ni	SCF/Bands	74	$68a_0$	17	US	PBE
Au	AIMD	36	0	32	NC	PZ

TABLE I: Key parameters of representative examples. K is the number of irreducible k-points; R_v is cell dimension in the vacuum direction; A is the number of atoms; PP is the pseudopotential type; and E_{xc} is the exchange-correlation functional.

the Fermi level to avoid constant shift errors. Agreement in force means the root mean square error is less than $10^{-3}Ry/au$ or 5% of the root mean square force, whichever is higher. Agreement in pressure means the absolute error is less than 1 kbar or the relative error is less than 5%. The total energy is an interesting metric mathematically, as it converges monotonically with respect to basis completeness due to variational freedom. However, it has no absolute physical meaning in pseudopotential calculations. Therefore, we report the total energy per atom (E/A), but do not constrain our methods to directly agree with respect to total energy.

E. PWDFT convergence

PWDFT has two primary discretizations: the fine k-point mesh and the number of plane waves, which is defined through a kinetic energy cut-off E_{cut}. Before comparing to the SRB, we ensure that that each structure is converged with respect to the number of k-points and plane-waves. The resulting configurations can be found in Table I.

For each structure, we converge with respect to the k-point mesh such that a doubling of the number of k-points on edge produces a result in agreement, as defined in Section IV D. For example, if the converged mesh had 16 points on edge, then a calculation with 32 points on edge should agree with it.

Convergence with respect to energy cutoff is a more nuanced problem, especially for ultra-soft pseudopotentials. For demonstration purposes, we simplify the matter by using a 32 (48) Rydberg cutoff for ultra-soft (norm conserving) pseudo-potentials in our presented examples. The validity of the SRB over a broad range of energy cut-offs is demonstrated in Figure 7.

The CNT and nickel slab structures are vacuum padded normal to the structure to model a reduced dimensional system in full periodic boundary conditions. At fixed atomic positions, the planar (axial) forces and stress of the Ni slab (CNT) are coupled to the amount of vacuum padding. This prevents agreement with respect to the raw parallel pressure and forces at reasonable vacuum separations. Instead, we allow each structure to relax normal to the vacuum direction before computing the agreement metrics. However, for comparison with the SRB result, we use the un-relaxed with the vacuum from the relaxed structure. This way, the forces and pressure are non-zero and can be meaningfully compared.

	N_Q	N_b	$\Delta E/A$	$\Delta\varepsilon$	ΔF	ΔP
Corners	8	1595	31.74	7.27	0.003617	173.26
Edges	20	3458	1.24	0.39	0.000231	1.29
Faces	26	3507	1.03	0.21	0.000092	0.55
Center	27	3509	1.00	0.20	0.000082	0.51

TABLE II: Comparison of basis size and accuracy for Au with various set of q-points. In all cases, $\sigma_b^2 = 10^{-7}$. Energies are given in meV, forces in Ry/a_0, and pressure in kbar.

σ_b^2	N_b	ΔE	$\Delta\varepsilon$	ΔF	ΔP
$10.^{-4}$	2173	99.62	13.640	0.002883	92.17
$10.^{-7}$	3507	1.03	0.209	0.000092	0.55
$10.^{-14}$	5485	0.50	0.043	0.000038	0.47
0.	5486	0.50	0.043	0.000038	0.47

TABLE III: Comparison of basis size and accuracy Au with various error tolerances, σ_b^2. In all cases, the q-points are corners, edges, and faces of the unit cube. Energies are given in meV, forces in Ry/a_0, and pressure in kbar.

Band structure computations are performed non-self consistently with the same energy cut-off and vacuum padding as the self-consistent calculation. However, in order to converge the band energies to high accuracy, sometimes the plane-wave eigensolver must be changed from Davidson to conjugate gradients. The switch is not uncommon for non-self consistent calculations.

V. PARAMETERIZATION

Thus far, we have tried to present the method as generally as possible. Here, we describe the relevant parameters of the method and their effects on the accuracy and performance of the SRB. In particular, we attempt to provide sensible defaults.

The times reported in this section should be interpreted qualitatively. For quantitatively valid timings, see Section VII.

A. Defining the SRB

There are three parameters that govern the subspace of the full basis in which the reduced basis is generated: the coarse sample of reciprocal space, Q, the truncation threshold, σ_b^2, and the maximum number of basis elements, N_{max}. Our implementation provides for a simultaneous definition of N_{max} and σ_b^2. As SCF calculations converge, the number of basis elements that satisfy σ_b^2

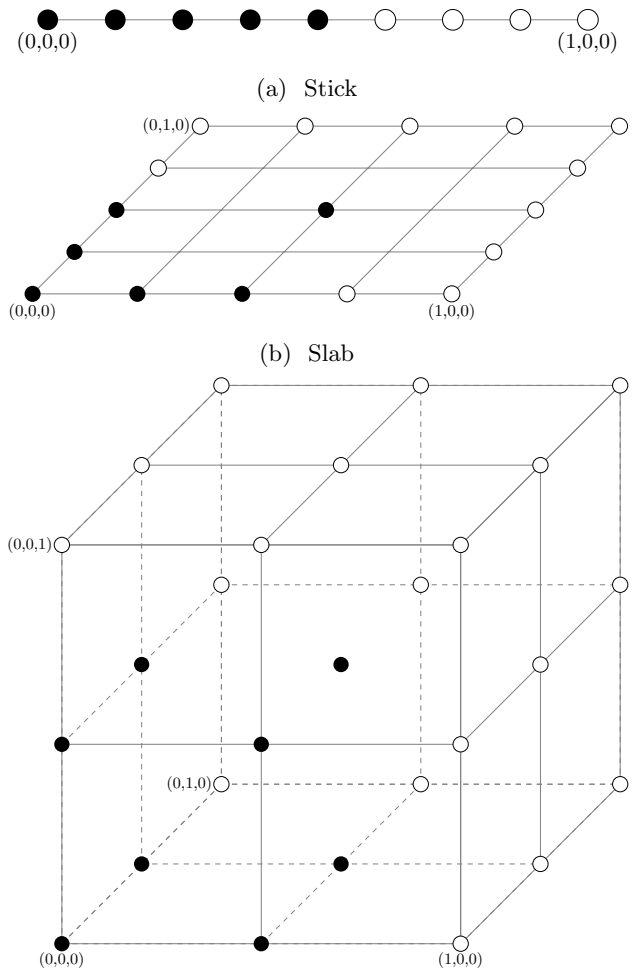

(a) Stick

(b) Slab

(c) Bulk.

FIG. 5: Coarse samples of the first Brillouin zone. Open points are related to closed points by inversion and translation.

tends to decrease, so if both are specified tightly, N_{max} applies early in the calculation and σ_b^2 later.

It was found that for non-self consistent calculations of sufficiently large unit cells, the Γ point and its mirrors on the unit cube are generally sufficient [10]. For self-consistent calculations, this is not generally the case. We find it best to use points that lie on the boundary of the K grid. For example, in three dimensions the corners, edge centers, and face centers of the unit cube should be used in that order. The relative basis sizes and accuracy for those coarse samples can be seem in Table II. In two dimensions, only the face $q_z = 0$ of the unit cube should be included. In one dimension, a 3-point sampling is generally not large enough to give the SRB enough degrees of freedom, so $(0, 1/8, 1/4, 3/8, 1/2, 5/8, 3/4, 1)$ is recommended. These coarse BZ samples can be seen in Figure 5.

On could consider increasing the number of input coarse eigenstates by artificially increasing the number of bands per coarse sample beyond the number desired in the fine sample. However, the additional high-energy eigenstates would do little to improve the SRBs representation of the occupied low-energy eigenstates. Simply, the high and low energy eigen-spaces overlap significantly less than the low energy eigenspaces across the BZ.

This is indicative of a more general trade-off: the more representative the coarse sample Q is of the fine sample K, the fewer basis elements are needed to achieve a fixed error tolerance. For a fixed threshold, selecting a more representative Q generally leads to a smaller basis. In addition to the aforementioned example of increasing the number of bands, this principal applies to applying symmetry beyond the first BZ.

In our experience, 5 to 15 basis elements are needed per band for the Q specified above. This corresponds to a threshold $\sigma_b^2 \approx 10^{-7}$.

σ_x^2	% Elem.	ΔE	$\Delta\varepsilon$	ΔF	ΔP
10^{-2}	10.24	51.59	73.75	0.000224	3.74
0.003	32.30	1.54	2.733	0.000094	0.53
10^{-3}	56.94	1.36	0.685	0.000093	0.54
10^{-4}	84.72	1.08	0.211	0.000092	0.55
10^{-6}	100.	1.03	0.209	0.000092	0.55
0	100.	1.03	0.209	0.000092	0.55

TABLE IV: Comparison of the number of auxillary basis elements, as a fraction of the total projector input states, and accuracy for Au with various auxillary error tolerances. In all cases, the auxillary q-point grid $Q = K$. Energies are given in meV, forces in Ry/a_0, and pressure in kbar.

System	Direct (s)	Auxillary basis (s)
CNT	47.98	20.59
Ni	3009.32	2612.15
Au	104.73	77.34

TABLE V: Comparison of the run-time of the direct transformation and auxiliary basis approaches to the projectors. In all cases, $\sigma_x^2 = 0.001$.

B. Auxiliary basis

The auxiliary basis $|x\rangle$ used in transforming the projectors is controlled by three parameters (Section III C) in the same manner as the SRB. Because the auxiliary basis need be constructed only once, but is used at every iteration, it is advantageous to be thorough in the sample Q. We let $Q = K$, which gives the variances of the basis elements statistical meaning and produces an 'optimal' basis. In this light, the threshold method is preferred over specifying the number of elements. A comparison of the auxiliary basis size and resulting accuracy for various thresholds is provided in Table IV. A

σ_ρ^2	% N_b	ΔE	$\Delta\varepsilon$	ΔF	ΔP
10^{-2}	23.4	6.42	11.93	0.000564	0.13
10^{-3}	29.6	8.07	3.79	0.000568	0.03
10^{-4}	37.0	8.11	2.636	0.000580	0.01
10^{-6}	48.2	8.11	2.635	0.000582	0.01
0	100.	8.11	2.620	0.000582	0.01
n/a	n/a	8.11	2.600	0.000537	0.01

TABLE VI: Comparison of the number of rank-1 components, as a fraction of the basis size, and accuracy for CNT with various density error tolerances. The last entry uses the direct transformation. Energies are given in meV, forces in Ry/a_0, and pressure in kbar.

System	Direct (s)	Density matrix (s)
CNT	511.75	47.68
Ni	2958.32	588.99
Au	4169.62	2223.18

TABLE VII: Comparison of the run time of the direct transformation and density approaches to computing the real-space electron density. In all cases, the density error tolerance $\sigma_\rho^2 = 10^{-4}$.

threshold $\sigma_x^2 = 10^{-3}$ is a robust choice.

The dimension of the SVD used to form the auxiliary basis scales linearly with the number of k-points. For small systems with very large number of k-points, this cost can dominate the overall cost of projection. In those cases, a smaller set of q-points, $Q \subset K$, should be used.

The construction of the auxiliary basis for each species is independent of the number of atoms of that species while the construction of the structure factors occurs once per atom. The overhead associated with constructing the auxiliary basis is amortized by the number of atoms. If the number of atoms of a species is small, the overhead costs are more of a factor. Therefore, it may only be beneficial to use the auxiliary basis when the number of atoms of a species is above some threshold, $N_{x,min}$.

Similarly, for non-self-consistent calculations there is only one projection. Thus, the overhead cost of the SVD is not amortized over many iterations. For non-self-consistent calculations, direct transformation of the projectors is recommended.

The three examples here are self-consistent with a moderate number of k-points. Thus, we expect the auxiliary basis to provide superior performance. It does, as seen in Table V.

C. Density matrix

There are two parameters that describe the building of the real-space density. The first is the threshold on the Fermi weight with which to add a wave function to the density matrix. This threshold is present in most PW DFT codes. The second is the threshold on the singular values of the density matrix. Because the density matrix is complete, this parameter provides an exact measure for the preservation of the electron density. A sensible value is 10^{-4}.

ϵ_b	N_I	Time (s)	ΔE	$\Delta\varepsilon$	ΔF	ΔP
10^{-2}	47	941.19	−1.7085	2.260	0.000851	−0.38
10^{-3}	47	1052.83	−1.6878	2.199	0.000932	−0.31
10^{-4}	49	1189.98	−1.6861	2.197	0.000944	−0.20
10^{-6}	50	1612.63	−1.6847	2.191	0.000947	−0.29
10^{-9}	50	2163.92	−1.6850	2.206	0.000949	−0.29
0	50	2347.25	−1.6878	2.199	0.000932	−0.31

TABLE VIII: Comparison of the number of iterations, run-time, and accuracy for Ni with various freeze thresholds. In all cases, the basis lifetime $n_L = 1$. Energies are given in meV, forces in Ry/a_0, and pressure in kbar.

n_L	N_I	Time (s)	ΔE	$\Delta\varepsilon$	ΔF	ΔP
5	52	1039.51	−3.879	2.308	0.000746	−0.35
4	47	1044.89	−3.879	2.303	0.000745	−0.35
3	56	1306.05	−3.879	2.299	0.000746	−0.34
2	54	1425.08	−3.879	2.300	0.000746	−0.34
1	52	1883.34	−3.879	2.296	0.000747	−0.33

TABLE IX: Comparison of the number of iterations, run-time, and accuracy for Ni with various basis life-times. In all cases, the freeze threshold $\epsilon_b = 10^{-6}$. Energies are given in meV, forces in Ry/a_0, and pressure in kbar.

D. Basis saving

There are two parameters that define the reuse of the SRB during self-consistent calculations: the basis lifetime, n_L, and the freeze threshold, ϵ_b. The basis lifetime is the minimum number of SCF iterations for which each basis is reused. The freeze threshold is the energy convergence criterion after which the basis is no longer updated.

As shown in Table IX, the accuracy is nearly completely independent of the basis lifetime. As the basis lifetime increases, the number of iterations to convergence, N_I, tends to increase. However, the mitigation of overhead costs generally makes up for the additional iterations, lowering overall runtime. Across the full range of systems, $n_L = 3$ is a robust choice.

As the wave functions and electron density converge, the SRB changes less and less between updates. When the energy eigenvalues are converged to ϵ_b, the standard convergence criteria for self-consistent calculations, the SRB is not updated. $\epsilon_b = 10^{-6}$ is a robust choice.

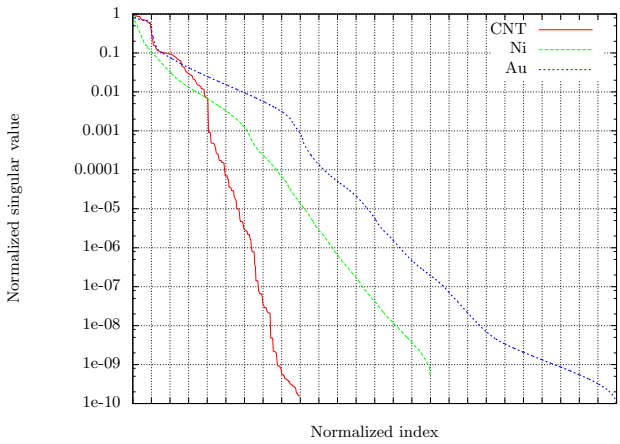

FIG. 6: Singular value spectrum taken from the end of a self-consistent calculation using the coarse samples specified in Figure 5. The vertical axis is the singular value divided by the number of coarse samples, that is the number of q-points Q. The horizontal axis is the singular value index divided by the number of eigenfunctions taken per q-point, that is the number of bands B.

VI. ACCURACY

There is an important inconsistency between the plane-wave basis used to construct the SRB and the plane-wave basis used in a plane-wave calculation. Plane-wave codes typically define their plane-wave bases as the set of g-vectors within a sphere centered around each k-point. This ensures that each k-point has the same well-defined E_{cut}, but each k-point therefore uses a different set of g-vectors.

$$G_{PW}(k) = \left\{g : |k+g|^2 < E_{cut}\right\} \tag{38}$$

The SRB, on the other hand, is k-independent; we typically define the SRB elements with respect to the Γ-point g-vectors. We could create a larger basis that includes every k-point basis by simply adding a g-vector in each direction. That is, a 28x34x60 reciprocal space grid would become a 30x36x62 grid. Similar issues arise when the mirroring procedure is applied to the input q-points [10], taking $q \to q'$. In the plane-wave code, the plane-wave basis at q and q' would be different.

In Quantum ESPRESSO, as in most modern PWDFT codes, a custom FFT interface is used which imposes an isotropic energy cut-off. We could increase the E_{cut} such that every g-vector at every plane-wave k-point is included in the Γ-point basis. Each SRB k-point would then be represented on a superset of g-vectors from a mix of plane-wave k-points.

$$G_{SRB} = \left\{g : |g|^2 < E_{cut} + |k_{max}|^2\right\} \supset \bigcup_{k\in K} G_{PW}(k) \tag{39}$$

If the original PW calculation is not fully converged with respect to the plane-wave cutoff, then the SRB with increased cut-off will converge to a different, but, in fact, more accurate result, due to their slightly larger variational freedom. However, the differences induced by inconsistent g-vector definitions closely track the g-vector truncation error. The error of the PW and SRB results, computed with respect to a fully g-vector converged result, are nearly identical until well below agreement standards defined in Section IV D. This effect is demonstrated in Figure 7. Generally, the convergence with respect to E_{cut} defines an agreement scale for the SRB result.

A. Convergence

The SRB spans a subspace of the span of the PW basis. Thus, as the SRB grows, it converges uniformly to the PW result. Furthermore, convergence of variational quantities, such as the total energy, is monotonic.

The basis elements are produced by singular value decomposition, Eq. 4. Singular value spectra of the three model systems can be seen in Figure 6. The steady exponential fall-off of the singular values demonstrates the redundancy of the set of input states $\{|u_{mq}\rangle\}$, for q sampled from Q, and suggests that the solution should converge exponentially with respect to the trace tolerance σ_b^2.

Convergence of numerical examples displays three regimes. First, when the SRB is too small, the nonlinearity of the system can cause a high-error plateau region. Subsequently, the error falls off exponentially. Finally, other differences between the cutoffs become apparent, causing the SRB difference to level off at the scale of the PW cutoff convergence.

There are other model differences related to the size of the auxiliary basis and the thresholds used for truncating sums in the density matrix. These errors are negligible when using the parameters described in Section V.

VII. PERFORMANCE

A simple empirical performance analysis is to break down the computation time into a setup and a per k-point cost, as seen in Table X. PWDFT spends the overwhelming majority of time in k-dependent terms, with only a few ρ-dependent operations. The SRB, on the other hand, spends a significant amount of time in k-independent 'overhead', in the form of the q-point calculations and basis transformations. The overhead is balanced by significantly lower per k-point costs. For the examples we have considered, the number of k-points needed to amortize the overhead, resulting in an overall faster calculation, ranges from 4 to 72 and The speedup in the limit of an infinite number of k-points ranges from 1.38 to 22x.

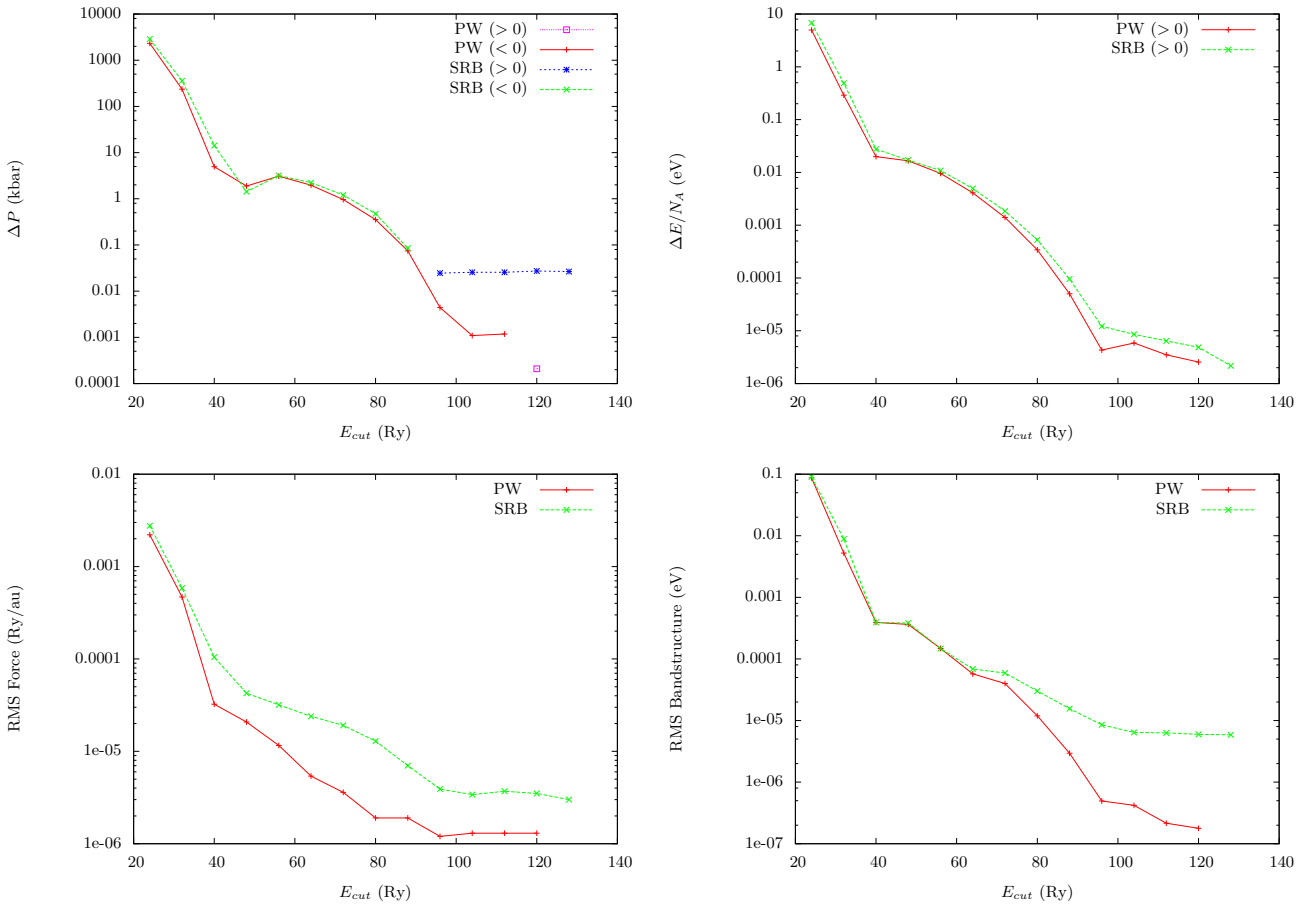

FIG. 7: Convergence of the total energy and pressure with respect to E_{cut} compared to $E_{cut} = 128$ Ry. The red (green) lines are PW (SRB) calculations. In the error in pressure is signed, so positive errors in the PW (SRB) calculation are in purple (red).

	PW-0	PW-K	PW Total	SRB-0	SRB-K	SRB Total	K_S	X_∞	X_{real}
CNT	196.06	91.13	3204.	470.25	3.99	602.	4	22.84	5.32
Ni	234.95	181.69	13680.	2551.6	52.28	6420.	18	3.47	2.13
Au	10.93	341.36	12300.	6783.55	246.57	15660.	72	1.38	0.78

TABLE X: Breakdown into k-independent and k-dependent costs. PW-K (SRB-K) is the time spent in k-depdenent steps of the PW (SRB) calculation divided by the number of k-points. PW-0 (SRB-0) is the time spent in k-indepdenent steps of the PW (SRB) calculation. K_S is the number of k-points needed for the SRB require less time than the PW calculation. X_∞ is the ratio of the PW to SRB calculation times in the limit as $K \to \infty$. X_{real} is the ratio of the PW to SRB calculation times for the actual calculation.

It should be noted that the results presented in Table X were optimized for minimal total time. The optimal parameters for these specific systems can be see in Table XIII. Greater per k-point acceleration could be achieved at the cost of greater overhead, and vice versa.

The SRB computation can be divided into the q-point plane-wave solutions, the construction of the SRB $\{|b\rangle\}$, the k-independent transformation of the Hamiltonian, the construction of the k-dependent Hamiltonian, diagonalization, and the construction of the electron density. The time spent in these sections of the code can be seen in Table XII.

For CNT, the SRB calculation is dominated by the plane-wave parts, which includes the q-point solutions and the k-independent transformation of the local potential. Diagonalization, on the other hand, takes less than 1% of the run-time. In this case the SRB can be thought of as a way to reduce the number of FFTs, which dominate the vacuum-padded calculation.

For Ni and Au, the focus shifts from plane-wave calculations to dense linear algebra. Au spends more time building the basis because bulk BZs have more q-points

	K	Q	G	S	α_D	α_S
CNT	33	5	79214	176	2.54	6.07
Ni	74	7	12431	972	2.02	11.85
Au	36	7	20672	3016	2.39	14.29

TABLE XI: Indicators of the relative performance of the SRB. K is the number of irreducible k-points; Q is the number of irreducible q-points; G is the number of plane waves; S is the number of SRB elements; α_D is the size of the subspace used in the Davidson algorithm as a ratio of the number of bands; and α_S is the size of the SRB as a ratio of the number of bands.

	PW	$\{\|b\rangle\}$	H	$H(k)$	Diag	ρ
CNT	62.6	4.7	19.9	6.4	0.7	5.8
Ni	16.0	2.6	12.8	17.2	43.8	7.7
Au	12.9	11.2	14.6	5.2	46.8	9.3

TABLE XII: Percent of run-time spent in plane-wave, basis construction, k-independent transformation, k-dependent construction, diagonalization, and density matrix sections, respectively.

than slab BZs. Ni spends more time constructing the k-dependent Hamiltonian and overlap matrix, S, because it employs ultra-soft pseudo-potentials. In both cases, diagonalization is a significant cost, but no individual cost dominates the calculation. The cost of each SRB step is at least linear in the number of basis elements, so reducing the basis size would reduce the cost of these calculations.

We can also reason about the performance of the SRB analytically. A full performance model, based on counting library calls, can be found in Appendix A. The model highlights three parameters as strongly indicative of the performance of the SRB compared to plane-waves: 1) the k-point ratio, Q/K; 2) the band density, B/V; and 3) the subspace inefficiency, α_S/α_D. The first two parameters are known a priori, while the third is buried deeply in the method. The value of these parameters for the three examples systems are found in Table XI. In all cases, lower values favor the SRB.

The k-point ratio, Q/K, characterizes the balance of over-head to per k-point costs. The lower the value, the less significant the overhead of the q-point plane-wave calculations. The k-point ratio captures the importance of the number of k-points on the relative efficiency of the SRB.

The band density provides an a priori estimate of the ratio of the SRB size to the number of plane-waves. Plane-waves cover space uniformly; the number of plane-waves is linearly dependent on the volume of the unit cell. The SRB is computed from the bands at each q-point. Thus, the size of the SRB depends linearly on the number of bands, not the volume. The ratio of the SRB size to the number of planes thus goes as the ratio of the number of bands to the volume, the band density. The band density captures the importance of the contents of the unit cell on the relative efficiency of the SRB.

The subspace inefficiency is defined as the ratio of the SRB size to the size of the iterative subspace that is directly solved in the iterative eigen-solver. Because the SRB currently uses a direct solver for diagonalization, this ratio measures the relative difficulty of direct diagonalization in the two methods. The cost of diagonalization is cubic in the dimension, so the relative costs are the cube of the subspace inefficiency. The subspace inefficiency captures the difficulty of spanning the entire BZ with a single basis, compared to computing an iterative subspace for each k-point individually. In that context, it is clear why the subspace inefficiency is greater than one.

Unlike the k-point ratio and band density, the subspace inefficiency can not be computed a priori. We have found it to depend strongly on the dimensionality of the BZ. The subspace inefficiency is highest for 3D BZ of bulk systems and lowest for the 1D BZ of nanowire, nanoribbons, and nanorods.

	σ_b^2	σ_x^2	σ_ρ^2	n_L	ϵ_b
CNT	$10.^{-6}$	0	$10.^{-4}$	1	$10.^{-2}$
Ni	$10.^{-5}$	$10.^{-3}$	$10.^{-4}$	5	$10.^{-5}$
Au	$10.^{-6}$	$10.^{-3}$	$10.^{-4}$	4	$10.^{-4}$

TABLE XIII: Optimal parameters for example systems such as to acheive agreement.

A. Parallel performance

The data presented thus far reflects serial performance as a proxy for model complexity. In reality, PWDFT problems are solved in parallel, motivating a presentation of parallel performance. However, evaluating the parallel performance of the SRB, particularly as compared to the full plane-wave method, is tricky.

Presently, Quantum ESPESSO provides a multi-tiered parallel framework with four key parameters: the number of k-points which are considered simultaneously, N_{pool}; the number of Fourier transforms which are evaluated simultaneously per k-point, N_{tg}; the number of processors used in parallel diagonalizations, N_{ort}; and the number of available processors, N_{proc}. Our SRB implementation has two parallel parameters: the number of k-points which are considered simultaneously, N_{pot}, and the number of available processors. N_{pot} can be thought of as N_{pool} for the reduced calculation while N_{pool} applies to the q-point calculation.

These models are inconsistent. For example, the PW implementation has no global context that spans pools. Operations on the density, ρ, are performed redundantly in each pool. This discourages the configuration $N_{pool} =$

	Default	Full	Relaxed	PW
CNT	949.	602.	544.	3204.
Au	29220.	15660.	8100.	12300.

TABLE XIV: Performance of the SRB based on serial run time (in seconds) using: default settings, settings which produce full agreement, and relaxed setting with reduced accuracy in agreement criteria specified in Section VII B.

N_{proc}, that is each k-point solved serially.

The SRB implementation does have a global context that spans pots. This context is necessary for the efficient construction of the SRB and k-independent transformations. It also applies to operations on the density. Therefore, for the SRB, $N_{pot} = N_{proc}$ is optimal, if there is sufficient memory available per process.

The optimal parallel configurations for the PW and SRB implementations are different. This makes direct comparison difficult.

B. Tuning

The performance results presented thus far come from parameter configurations that satisfy all three agreement conditions from Section IV D. In some cases, these agreement conditions can be relaxed. For example, if the geometry is prescribed, then force and stress agreement for the purpose of relaxing the structure is unnecessary. On the other hand, if the end result is a molecular dynamics trajectory, then a highly accurate band structure may not be required at every time-step. The SRB provides a means for expressing lessened accuracy expectations to yield higher throughput calculations.

Here, we explore two situations with reduced accuracy requirements. The first is AIMD, where the band-structure and stress at every time step are not important. We further relax the force constraint to be within 10% of the root mean square magnitude. The second is a band structure calculation of a prescribed CNT, where the forces and stress are not important. The constraint on the root mean square error of the band structure is kept at 5 meV.

The relaxed accuracy constraints lead to improved performance, seen in Table XIV. The resulting SRB performance for is 1.52x better than plane-waves for Au and 5.89x better for CNT.

VIII. DISCUSSION

A. Where to expect speedup

The performance model is able to capture the major indicators of SRB performance: the ratio of k-points to q-points, the ratio of plane-waves to reduced basis elements, and the ratio of reduced basis elements to bands.

Three dimensional bulk materials in modest cells often require a large number of k-points to resolve the Fermi surface. When disordered, broken symmetries further increase the number of irreducible points. However, bulk materials also have relatively high band densities. Furthermore, the entire volume of the BZ must be covered by the basis, requiring more basis elements per band than a 1D or 2D BZ. Using direct diagonalization, as is presently done, the SRB is unlikely to provide a significant reduction in the run-time. Using iterative diagonalization, one could reduce the performance model's dependence on the ratio of basis elements to bands, providing a speedup through more efficient application of the Hamiltonian.

One dimensional wire, tube, or ribbon materials require vacuum padding to imitate isolation in periodic boundary conditions. The padding increases the number of plane-waves without changing the number of occupied bands required to build the electron density. The BZ is also one-dimensional, so relatively few basis elements per band are needed. However, a linear BZ generally means fewer k-points. One dimensional systems frequently approach the limit given by Equation A6 with $Q \approx 4$ and $\alpha_S \approx 2\alpha_D$.

Slab geometries are the best of both worlds. The need for vacuum padding increases the number of plane-waves compared to the bulk case. The two planar dimensions can require many k-points compared to the one-dimensional case. The k-points lie in a plane, which can be spanned my a smaller number of basis functions than the bulk.

B. Implications for parallelism

The SRB adds a global operation to the SCF process: the construction of the basis. However, the majority of the effort in building the basis comes from constructing the overlap matrix, Eq. 3, and the expansion of the reduced basis elements into the plane-wave transformation matrix, Eq 6. These operations can be easily parallelized. Only the diagonalization of the covariance matrix, Eq. 4, exhibits poor scaling, and that operation is both G and K independent.

A significant advantage of the SRB comes in the reduction of the problem size. The reduced, dense eigenvalue problem is significantly smaller than the plane-wave equivalent. For example, a system with 1000 bands and 20 basis elements per band would fit on a 16 GiB node. Generally, the SRB requires less memory and therefore fewer nodes to handle each k-point. This allows for aggressive k-point parallelism, which further improves scalability by keeping the problem local.

The operations required to build the k-independent parts of the Hamiltonian require more memory. The transformation matrix, $\langle b|g\rangle$, can be particularly large. However, the application of these matrices is through

highly parallel matrix-matrix products. Therefore, a scheme which parallelizes the construction of a set of dense Hamiltonians over a set of nodes and then subdivides into maximally local solves parallelizes very well compared to the plane-wave counterpart.

Additionally, the SRB uses significantly fewer FFT operations than the plane-wave counterpart. Parallel FFTs have high communication to computation ratios and are known to limit plane-wave DFT scalability in many cases. Avoiding the majority of them lessens this parallel bottleneck.

C. Implications for accelerators

Accelerator systems, including those based on GPUs and coprocessors, can struggle to provide sufficient network bandwidth to keep up with their high floating-point and local memory performance. The aforementioned decrease in problem size, increase in possible locality, and avoidance of parallel FFTs reduces the communication overhead significantly. The dense linear algebra that comprises the majority the SRB computation is an ideal case of accelerators. Early testing indicates that the potential for the application of accelerators to the SRB exceeds that of traditional PW calculations.

D. Advantages for norm-conserving pseduo-potentials

The SRB presents three advantages for norm-conserving pseudo-potentials. First, the NCPP non-local potential is static. When the SRB is saved, only the local potential must be recomputed between SCF iterations. Second, the eigensystem does not need to be generalized which greatly improves the efficiency of direct solvers compared to subspace methods that still require generalized solves. Third, the usual disadvantage of NCPPs is the requirement of more plane-waves than USPPs. The SRB is relatively insensitive to the number of plane-waves compared to full plane-wave calculations. We can expect the SRB to reduce the performance gap between NCPPs and USPPs.

E. Comparison to NSCF calculations

There are three primary differences between self-consistent and non-self-consistent calculations in the SRB. The first, and most obvious, is the addition of a step to construct the charge density. The method outlined in Section III E is always cheaper than the transformation of the local potential, so this should not significantly impact the performance of the method.

The second difference is the reduction in the number of k-points. The density does not depend as strongly on the k-point sampling as the band-structure. It is not uncommon for NSCF calculations to exceed thousands of k-points. It is rare for self-consistent calculations to reach one hundred. This shifts the target away from creating the smallest, most expressive basis towards methods that facilitate cheap transformations, reducing the overhead cost.

The third difference is related to the diagonalization. Iterative diagonalization schemes take as input a convergence threshold for the accuracy of the eigenvalues. The first self-consistent iteration specifies a large threshold. Subsequent iterations use successively smaller thresholds until the final criteria is met. Because each SCF iteration uses the previous iterations result as an 'initial guess', the diagonalizer has very little work to do at each iteration. In a sense, self-consistent calculations only diagonalize once, but they interleave that diagonalization with self-consistent updates to the Hamiltonian. In NSCF calculations, the same iterative diagonalizer is used but the convergence threshold can not be reduced. Therefore, NSCF diagonalization takes much longer than a single SCF iteration.

The SRB uses direct diagonalization of the dense Hamiltonian, which doesn't benefit much from a reduced convergence threshold. Thus, the reduction in diagonalization time is much more pronounced for NSCF calculations, even if the number of k-points, plane-waves, and basis elements were identical.

IX. FUTURE WORK

A. Preconditioned iterative diagonalization

We have seen the direct diagonalization step to be a significant bottleneck for higher dimensional systems. The poor performance is rooted in the larger size of the SRB compared to the subspace formed in the Davidson iteration employed by the plane wave code. Direct diagonalization is cubic in the matrix dimension, so even a doubling of the SRB size compared to the subspace will degrade performance by 8x.

The obvious solution is to use an iterative method, such as Davidson, to diagonalize the SRB Hamiltonian. This would allow the SRB to take advantage of gradual diagonalization inherit to self-consistent calculations. The iterative subspace needed to diagonalize the SRB Hamiltonian should be no larger than the PWDFT counterpart, and could indeed be smaller.

The challenge here, as pointed out by Shirley [6], is to pick a preconditioner. The SRB Hamiltonian is dense and need not be diagonally dominant. The SRB does provide a factorization of the Hamiltonian with respect to k. One could consider directly inverting the k-independent part of the Hamiltonian and re-using it as the foundation of dense preconditioner. Such a preconditioner could be much more accurate than the diagonal preconditioners used in the plane-wave code, enabled by the small dense matrix size.

B. Pruning

Taking temporal redundancy a step farther, one could not only reuse a previous basis but reduce it further based on statistics of the previous iteration. For example, on the first iteration, the reduced basis could be produced with no truncation. Then, at the end of the iteration, the weight of the states $\langle b_i|u_{nk}\rangle$ on each basis element would indicate which elements to remove. This would amount to computing scores

$$\omega_i = \sum_{nk} f(\epsilon_{nk}) \langle b_i|u_{nk}\rangle \tag{40}$$

and removing some of the elements with the lowest score. This could be up to a sum of missing scores or a ratio of the number of elements. We call this pruning. Pruning could be repeated at the end of each iteration until the basis is frozen or a new, un-truncated basis could be formed.

There are two advantages of pruning over simply producing a new basis. First, pruning takes into consideration the Bloch states at all the k-points, not just the q-points. Therefore, we could expect pruning to better represent the basis elements which are being used. Secondly, pruning does not require a new transformation of the density-independent components of the Hamiltonian. Instead, the pruned Hamiltonian is formed by eliminating rows and columns from SRB Hamiltonian and rows from the set of projectors.

C. Dividing the BZ

In the present scheme, the entire first BZ must be covered by a single reduced basis. This requirement is the primary contributor to subspace inefficiency. In Davidson, the subspace need only be valid for a single k-point.

The BZ could be divided into regions, each covered by a smaller basis. For example, the first BZ could be split into octants, the first octant covering $[0, 1/2]^3$ in crystal coordinates. Dividing the BZ in this way would increase the overhead costs, but reduce the per k-point cost. If the number of k-points were very large, the reduction in per-k-point cost due to the smaller reduced bases could outweigh the added overhead cost.

D. Uncertainty quantification

To get an even more rigorous bound on the error in basis truncation *and subspacing*, one could intermittently solve the problem in the plane-wave basis on the fine mesh and compute $\langle B_i|\bar{\psi}_{nq}\rangle$. The compact basis could then be constructed with far more statistical information.

Taking it a step farther, the choice of the coarse mesh itself could be informed by statistical analysis of the redundancy in full-basis solutions on the fine mesh. This would allow for a scheme in which a single error tolerance is specified and the algorithm statistically conforms to that tolerance.

E. Exact exchange

The computation of the exact-exchange operator in the reduced basis presents many challenges. The exchange operator can be written as

$$\langle b_i\,|K(k)|\,b_j\rangle = \sum_{nk} f_{nk}(\varepsilon) \frac{\langle b_i|r\rangle \langle r|\psi_{nk}\rangle \langle \psi_{nk}|r'\rangle \langle r'|b_j\rangle}{|r-r'|}, \tag{41}$$

up to constant prefactors. This expression is generally treated by considering a two-particle density:

$$\rho_{a,b}(r) = \langle \phi_a|r\rangle \langle r|\phi_b\rangle\,, \tag{42}$$

where ϕ_a and ϕ_b are arbitrary complex fields (e.g. the wave function ψ_{nk} or basis element b_j). Then an analogous two-particle potential is formed through a Poisson solve. Both of these steps are defined in specific bases, real space for the density and reciprocal space for the Poisson solve. As linear operators, the two operators are a 2,2 tensor, which is too large to be direct transformed or stored in the SRB.

Another approach would be to recognize that the sum over occupied states produces a non-diagonal density operator:

$$\langle r\,|\rho|\,r'\rangle = \sum_{nk} f_{nk}(\varepsilon) \langle r|\psi_{nk}\rangle \langle \psi_{nk}|r'\rangle\,. \tag{43}$$

We would like to perform a rank-1 decomposition, as in Section III E, but the off-diagonal elements of the density do benefit from cancellation of the k-point phase. That is, the non-diagonal density operator is full rank.

However, it is not the opinion of the authors that the exact exchange operator can not be expressed in the SRB. Simply, this is an area which requires further study.

X. CONCLUSIONS

We have shown the SRB to be capable of achieving an arbitrary degree of accuracy compared to the host calculation, in this case plane-waves. More generally, the SRB completes to the primary basis and therefore inherits the accuracy and convergence properties of that basis. For plane-waves, this makes the SRB uniformly convergent.

The current implementation of the SRB is more efficient than plane-wave calculations for reduced dimensional extended systems, such as surfaces, nanorods, nanowires, and nanoribbons. The reduction in runtime depends chiefly on the number of k-points, the band density, and the dimensionality of the BZ. We have demonstrated speed ups in highly accurate self-consistent calculations in excess of 5x. Reducing accuracy constraints

to be more in line with practical expectations for MD trajectories leads to a 1.52x performance improvement for bulk Au.

To enable highly efficient self-consistent calculations, we have described a k-independent transformation of the PWDFT Hamiltonian into the SRB and analogous transformation of the electron density to real-space. The k-independent transformation relies on the novel auxiliary basis approach for the projectors and explicit density matrix approach for the electron density. These techniques are examples of a broader class of k-independent transformations that could be used to efficiently transform other operators into the SRB. For example, the treatment of electron-phonon interactions in density functional perturbation theory[37] and the treatment of electronic excited states within many-body perturbation theory [38, 39], in some systems, lead to large numbers of matrix elements that need to be computed on a fine k-point grid.

The SRB shifts numerical focus away from FFTs and towards dense linear algebra. For example, the number of FFTs is independent of the number of k-points. Dense linear algebra has a very different computational profile than FFTs. Particularly, dense linear algebra is generally floating-point bound, while FFTs are bandwidth bound. With the growing heterogeneity of computer architectures, this difference will have growing implications for the performance of the methods. If general purpose graphical processing units and other modern coprocessors are any indication, floating-point performance will continue to grow more rapidly than bandwidth, favoring dense linear algebra.

We have outlined where there is room for improvement in our current implementation of the SRB. The most pressing issue is iterative diagonalization, which will greatly improve the performance of 2D and 3D systems, as demonstrated by a Ni slab and a finite-temperature Au snapshot here. Additionally, the coarse meshes, q-points, recommended here are independent of the composition of the fine mesh, k-points, which is likely sub-optimal. Pruning could provide an efficient means for updating the basis without performing additional plane-wave calculations or transformations. When the number of k-points is large, likely in non-self-consistent calculations, dividing the BZ could reduce the subspace inefficiency. Therefore, at this early stage, with room for development in theory and in implementation, we deem the SRB approach to hold much promise for improving the efficiency of self-consistent DFT calculations.

ACKNOWLEDGMENTS

M Hutchinson acknowledges support from the Department of Energy Computational Science Graduate Fellowship Program (Grant No. DE-FG02-97ER25308). This work was performed at the Molecular Foundry, supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Appendix A: Performance model

Both PWDFT and SRB calculations spend the majority of their runtime in the FFT, BLAS, and LAPACK libraries. The performance of the two methods can be well characterized by the number and size of those calls.

Consider a self-consistent calculation of an electron density using a traditional plane-wave method. At the heart of the procedure is a matrix free iterative eigensolver, most commonly conjugate gradient (CG) [40] or Davidson [11, 41]. Both access the Hamiltonian by computing the action on a vector $H|\phi\rangle$. We restrict the rest of this discussion to Davidson, which is considered higher performance and a good representative of related sub-space techniques. Davidson uses the action to represent the Hamiltonian and overlap matrix in a dense subspace [42]. In PWDFT, each Hamiltonian evaluation requires 2 fast Fourier transforms (FFT) to access the real-space local potential, and 2 matrix-vector (MV) products to project into and out of the atomic wave function space. The dense Hamiltonian is formed by inner-products over plane-waves. The subspace eigenproblem is then solved directly. Candidate solutions are transformed back into the full space and their residuals, $(H-\varepsilon S)|\psi\rangle$, are used to expand the subspace using a diagonal approximation of the inverse. The construction of the electron density, ρ, requires an additional FFT and MV product per eigenfunction. In sum, the Davidson workload per k-point per self-consistent iteration is:

$$W_{PW} = (2\alpha_D + 1)B\left(\mathrm{FFT}[G] + \mathrm{MV}[P,G]\right) + MM[\alpha_D B, \alpha_D B, G] + \mathrm{EVD}[\alpha_D B] \quad \text{(A1)}$$

where B is the number of bands, G is the number of plane-waves, P is the number of projectors, α_D is size of the final subspace as multiple of the number of bands, and $\mathrm{EVD}[N]$ denotes an eigenvalue decomposition (eigensolve) of size N which computes all eigenvectors. The assumption is made that the direct diagonalization of the largest subspace dominates the other diagonalizations, which underestimates the subspace diagonalization costs. Note that in Davidson, the matrix-vector (MV) products can be coalesced into a more efficient matrix-matrix (MM) product because the subspace can be formed in parallel. In some codes, the transformation to pseudo-wave function space is performed in real-space, in which case the MV product is sparse with a number of non-zero terms that grows linearly with the system size, not quadratically.

When using the SRB, as described in Section III, the

workload becomes:

$$\begin{aligned} W_{SRB} = &\ \mathrm{SVD}[\bar{Q}B, G] + \mathrm{MM}[\bar{Q}B, \bar{Q}B, G] + \mathrm{MM}[\bar{Q}B, S, G] \\ &+ 2S\ \mathrm{FFT}[G] + 5\mathrm{MM}[S, S, G] \\ &+ K\ \mathrm{MM}[S, P, X] + A\ \mathrm{MM}[S, X, G] \\ &+ K\ \mathrm{MM}[S, S, P] \\ &+ K\ \mathrm{EVD}[S] \\ &+ K\ \mathrm{MM}[S, S, B] + \mathrm{SVD}[S, S] \\ &+ \alpha_\rho S\left(\mathrm{FFT}[G] + \mathrm{MV}[S, G]\right) \\ &+ (2\alpha_D)QB\left(\mathrm{FFT}[G] + \mathrm{MV}[P, G]\right) \end{aligned} \tag{A2}$$

Where SVD denotes the singular value decomposition, S is the size of the SRB, X is the size of the auxiliary basis, A is the number of atoms, $\bar{Q}$ is the number of q-points, Q is the number of irreducible q-points, EVD denotes the eigenvalue decomposition, and α_ρ is the fraction of the SRB degrees of freedom needed to represent the density matrix, with a typical value of 1/4.

When a saved basis is used, the workload reduces to:

$$\begin{aligned} W_{SRB'} = &\ S\mathrm{FFT}[G] + \mathrm{MM}[S, S, G] \\ &+ K\mathrm{EVD}[S] \\ &+ K\mathrm{MM}[S, S, B] + \mathrm{SVD}[S, S] \\ &+ \alpha_\rho S\left(\mathrm{FFT}[G] + \mathrm{MV}[S, G]\right) \end{aligned} \tag{A3}$$

$$\tag{A4}$$

which omits the plane-wave calculation and SVD used to form the basis and the transformation of projectors and kinetic energy, which are density independent, and uses the previously computed Fourier transform of the basis elements.

As the number of k-points increases, the cost of the EVD dominates all other SRB costs. The Davidson solver in the plane-wave algorithm includes a similarly sized EVD per k-point. We can write $S = \alpha_S B$ so that the ratio α_D/α_S characterizes the relative sizes of the direct eigensolves. In this large K regime, the speedup should be:

$$\mathrm{Speedup} = \left(\frac{\alpha_D}{\alpha_S}\right)^3 \left[1 + \frac{W_{PW} - \mathrm{EVD}[\alpha_D B]}{\mathrm{EVD}[\alpha_D B]}\right] \tag{A5}$$

When $\alpha_D/\alpha_S < 1/2$, the performance of the reduced diagonalization degrades rapidly. In these cases, $\alpha_S >> 1$, so most of the direct solver's effort is wasted on uninteresting eigenpairs. An iterative eigensolver should be used instead.

As the number of plane-waves per SRB matrix element, G/S, increases, the cost of the G-dependent terms dominates all other costs. The Fourier transforms are the highest order in G, so we only count them. The plane-wave calculation uses $2KB\alpha_D$ while the SRB uses $2QB\alpha_D + 2S = 2B(Q\alpha_D + \alpha_S)$. The speed-up should approach

$$\mathrm{Speedup} = \frac{K\alpha_D}{Q\alpha_D + \alpha_S} \tag{A6}$$

Because G grows linearly with the volume but S grows linearly with the number of bands, B, the ratio $S/G \approx B/V$, the *band density*. The lower the band density, the better one can expect the SRB to perform.

[1] A. Jain, S. P. Ong, G. Hautier, W. Chen, W. D. Richards, S. Dacek, S. Cholia, D. Gunter, D. Skinner, G. Ceder, and K. A. Persson, APL Materials **1**, 1 (2013).
[2] W. Kohn and L. Sham, Physical Review **385** (1965).
[3] A. Belsky, M. Hellenbrandt, V. L. Karen, and P. Luksch, Acta Crystallographica Section B: Structural Science **58**, 364 (2002).
[4] Z. Wang, Y. Sun, X.-Q. Chen, C. Franchini, G. Xu, H. Weng, X. Dai, and Z. Fang, Physical Review B **85**, 195320 (2012).
[5] R. Martin, *Electronic structure: basic theory and practical methods* (Cambridge Univ Pr, 2004).
[6] E. Shirley, Physical Review B **54**, 464 (1996).
[7] J. Luttinger and W. Kohn, Physical Review **173** (1955).
[8] G. Dresselhaus, A. Kip, and C. Kittel, Physical Review **98** (1955).
[9] M. Rathinam and L. R. Petzold, SIAM Journal on Numerical Analysis **41**, 1893 (2003).
[10] D. Prendergast and S. G. Louie, Physical Review B **80**, 1 (2009).
[11] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. Dal Corso, S. de Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougoussis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri, R. Mazzarello, S. Paolini, A. Pasquarello, L. Paulatto, C. Sbraccia, S. Scandolo, G. Sclauzero, A. P. Seitsonen, A. Smogunov, P. Umari, and R. M. Wentzcovitch, Journal of Physics: Condensed Matter **21**, 395502 (2009).
[12] `http://www.github.com/maxhutch/qe-srb`.
[13] Pickard and R. Needs, Physical Review Letters **107**, 1 (2011).
[14] K. T. Chan, J. B. Neaton, and M. L. Cohen, Physical Review B **77**, 235430 (2008).
[15] D. J. Lucia, P. S. Beran, and W. a. Silva, Progress in Aerospace Sciences **40**, 51 (2004).
[16] L. Sirovich, Quarterly of Applied Mathematics **XLV**, 561 (1987).
[17] C. Ding and X. He, in *Proceedings of the twenty-first international conference on Machine learning* (2004).
[18] C. Homescu, L. Petzold, and R. Serban, SIAM Journal on Numerical Analysis **43**, 1693 (2005).
[19] B. Epureanu, Journal of Fluids and Structures **17**, 971 (2003).
[20] P. Umari, G. Stenuit, and S. Baroni, Physical Review B **79**, 7 (2009).
[21] C. Persson and C. Ambrosch-Draxl, Computer Physics Communications **177**, 280 (2007).
[22] M. I. Hussein, Proceedings of the Royal Society A: Mathematical, Ph
[23] G. Pau, Physical Review E **76**, 046704 (2007).
[24] D. Vanderbilt, Physical Review B **25**, 4228 (1990).

[25] P. Blöchl, Physical Review B **41**, 5414 (1990).
[26] P. Blöchl, Physical Review B **50**, 17953 (1994).
[27] W. Pickett, Computer Physics Reports (1989).
[28] L. Kleinman and D. Bylander, Physical Review Letters **48**, 1425 (1982).
[29] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammerling, and A. McKenney, *LAPACK Users' guide* (Siam, 1999).
[30] K. Laasonen, A. Pasquarello, R. Car, C. Lee, and D. Vanderbilt, Physical Review B **47**, 142 (1993).
[31] O. Nielsen and R. Martin, Physical Review Letters **50**, 697 (1983).
[32] `https://github.com/maxhutch/srb-supplemental`.
[33] D. Krepel and O. Hod, Surface Science **605**, 1633 (2011).
[34] P. Ganesh and M. Widom, Journal of Non-Crystalline Solids **357**, 442 (2011).
[35] X.-Z. Li, M. I. J. Probert, A. Alavi, and A. Michaelides, Physical Review Letters **104**, 066102 (2010).
[36] Y. Ping, D. Hanson, I. Koslow, T. Ogitsu, D. Prendergast, E. Schwegler, G. Collins, and a. Ng, Physics of Plasmas **15**, 056303 (2008).
[37] F. Giustino, M. Cohen, and S. Louie, Physical Review B **76**, 165108 (2007).
[38] M. Hybertsen and S. Louie, Physical Review B **34** (1986).
[39] M. Rohlfing and S. Louie, Physical Review B **62**, 4927 (2000).
[40] M. Payne, M. Teter, and D. Allan, Reviews of Modern Physics **64**, 1045 (1992).
[41] G. Kresse, Computational Materials Science **6**, 15 (1996).
[42] E. Davidson, Journal of Computational Physics , 87 (1975).